Express Mail Label No. EL241884235US

A

# UTILITY PATENT APPLICATION TRANSMITTAL
## (Large Entity)

*(Only for new nonprovisional applications under 37 CFR 1.53(b))*

Docket No.
12463 (CA998-012)

Total Pages in this Submission
3

jc672 U.S. PTO 06/10/99

jc510 U.S. PTO 09/329558 06/10/99

**TO THE ASSISTANT COMMISSIONER FOR PATENTS**
**Box Patent Application**
**Washington, D.C. 20231**

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

**MAPPING A STACK IN A STACK MACHINE ENVIRONMENT**

and invented by:

**Graham Chapman** **Andrew Low**
**John Duimovich**
**Trent Gray-Donald**
**Graeme Johnson**

**If a CONTINUATION APPLICATION,** *check appropriate box and supply the requisite information:*

☐ **Continuation** ☐ **Divisional** ☐ **Continuation-in-part (CIP)** of prior application No.: _____

Which is a:

☐ **Continuation** ☐ **Divisional** ☐ **Continuation-in-part (CIP)** of prior application No.: _____

Which is a:

☐ **Continuation** ☐ **Divisional** ☐ **Continuation-in-part (CIP)** of prior application No.: _____

Enclosed are:

**Application Elements**

1. ☒ Filing fee as calculated and transmitted as described below

2. ☒ Specification having **thirty (30)** pages and including the following:

   a. ☒ Descriptive Title of the Invention
   b. ☐ Cross References to Related Applications *(if applicable)*
   c. ☐ Statement Regarding Federally-sponsored Research/Development *(if applicable)*
   d. ☐ Reference to Microfiche Appendix *(if applicable)*
   e. ☒ Background of the Invention
   f. ☒ Brief Summary of the Invention
   g. ☒ Brief Description of the Drawings *(if drawings filed)*
   h. ☒ Detailed Description
   i. ☒ Claim(s) as Classified Below
   j. ☒ Abstract of the Disclosure

# UTILITY PATENT APPLICATION TRANSMITTAL
## (Large Entity)

*(Only for new nonprovisional applications under 37 CFR 1.53(b))*

| Docket No. |
|---|
| 12463 (CA998-012) |
| Total Pages in this Submission |
| 3 |

**Application Elements (Continued)**

3. ☒ Drawing(s) *(when necessary as prescribed by 35 USC 113)*
   a. ☒ Formal Number of Sheets **fourteen (14)**
   b. ☐ Informal Number of Sheets _____
4. ☒ Oath or Declaration
   a. ☒ Newly executed *(original or copy)* ☐ Unexecuted
   b. ☐ Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional application only)*
   c. ☒ With Power of Attorney ☐ Without Power of Attorney
   d. ☐ *DELETION OF INVENTOR(S)*
   Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference *(usable if Box 4b is checked)*
   The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Computer Program in Microfiche *(Appendix)*
7. ☐ Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all must be included)*
   a. ☐ Paper Copy
   b. ☐ Computer Readable Copy *(identical to computer copy)*
   c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

**Accompanying Application Parts**

8. ☒ Assignment Papers *(cover sheet & document(s))*
9. ☐ 37 CFR 3.73(B) Statement *(when there is an assignee)*
10. ☐ English Translation Document *(if applicable)*
11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☒ Certificate of Mailing
    ☐ First Class ☒ Express Mail *(Specify Label No.):* **EL241884235US**

# UTILITY PATENT APPLICATION TRANSMITTAL
## (Large Entity)

*(Only for new nonprovisional applications under 37 CFR 1.53(b))*

| Docket No. |
|---|
| 12463 (CA998-012) |
| Total Pages in this Submission |
| 3 |

**Accompanying Application Parts (Continued)**

15. ☒ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. ☒ Additional Enclosures *(please identify below):*

**Claim of Priority**
**Executed Associate Power of Attorney and Request for Change of Mailing Address**

**Fee Calculation and Transmittal**

**CLAIMS AS FILED**

| For | #Filed | #Allowed | #Extra | Rate | Fee |
|---|---|---|---|---|---|
| Total Claims | 23 | - 20 = | 3 | x $18.00 | $54.00 |
| Indep. Claims | 3 | - 3 = | 0 | x $78.00 | $0.00 |
| Multiple Dependent Claims (check if applicable) ☐ | | | | | $0.00 |
| | | | | BASIC FEE | $760.00 |
| OTHER FEE *(specify purpose)* | Recordation of Assignment | | | | $40.00 |
| | | | | TOTAL FILING FEE | $854.00 |

☐ A check in the amount of to cover the filing fee is enclosed.

☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. **50-0510/IBM** as described below. A duplicate copy of this sheet is enclosed.

- ☒ Charge the amount of **$854.00** as filing fee.
- ☒ Credit any overpayment.
- ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

*Signature*

Dated: **June 10, 1999**

**Richard L. Catania, Reg. No.: 32,608**
**Scully, Scott, Murphy & Presser**
**400 Garden City Plaza**
**Garden City, New York 11530**
**(516) 742-4343**

cc:

**PATENT**

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

**Applicant:** Graham Chapman, et al. **Docket:** 12463

**Serial No.:** To be assigned **Dated:** May 31, 1999

**Filed:** Herewith

**For:** MAPPING A STACK IN A STACK MACHINE ENVIRONMENT

Assistant Commissioner for Patents
Washington, D.C. 20231

**ASSOCIATE POWER OF ATTORNEY AND REQUEST FOR CHANGE OF MAILING ADDRESS**

Sir:

Applicant(s), by (his/her/their) attorneys of record, hereby grant(s) an Associate Power of Attorney to:

RICHARD L. CATANIA, Reg. No. 32,608; FRANK S. DIGIGLIO, Reg. 31,346; KENNETH L. KING, Reg. No. 24,223; STEPHEN D. MURPHY, Reg. No. 22,002; LEOPOLD PRESSER, Reg. No. 19,827; and JOHN S. SENSNY, Reg. No. 28,757

with full power of substitution to prosecute this application and transact all business in the United States Patent and Trademark Office in connection therewith.

Applicant(s) further request(s) that all future correspondence in connection with this application be directed and addressed to:

RICHARD L. CATANIA, ESQ.
SCULLY, SCOTT, MURPHY AND PRESSER
400 Garden City Plaza
Garden City, New York 11530
Direct all telephone calls to: (516) 742-4343.

Respectfully submitted

Jay P. Sbrollini
Reg. No. 36,266

International Business Machines Corporation
Thomas J. Watson Research Center
Route 134/Kitchawan Road
P.O. Box 218
Yorktown Heights, New York 10598

f:\work\105\asspoa.wpd

# MAPPING A STACK IN A STACK MACHINE ENVIRONMENT

## BACKGROUND OF THE INVENTION

### Technical Field

This invention relates generally to the field of memory optimization, and provides, in particular, a method for mapping the dynamic memory stack in a programming language environment such as Java.

### Prior Art

Java programs (as well as those in other object-oriented or OO languages) require the allocation of dynamic storage from the operating system at run-time. This run-time storage is allocated as two separate areas known as the "heap" and the "stack". The stack is an area of addressable or dynamic memory used during program execution for allocating current data objects and information. Thus, references to data objects and information associated with only one activation within the program are allocated to the stack for the life of the particular activation, Objects (such as classes) containing data that could be accessed over more than one activation must be heap allocated or statically stored for the duration of use during run-time.

Because modern operating systems and hardware platforms make available increasingly large stacks, modern applications have correspondingly grown in size and complexity to take advantage of this available memory. Most applications today use a great deal of dynamic memory. Features such as multitasking and multithreading

increase the demands on memory. OO programming languages use dynamic memory much more heavily than comparable serial programming languages like C, often for small, short-lived allocations.

The effective management of dynamic memory, to locate useable free blocks and to deallocate blocks no longer needed in an executing program, has become an important programing consideration. A number of interpreted OO programming languages such as Smalltalk, Java and Lisp employ an implicit form of memory management, often referred to as garbage collection, to designate memory as "free" when it is no longer needed for its current allocation.

Serious problems can arise if garbage collection of an allocated block occurs prematurely. For example, if a garbage collection occurs during processing, there would be no reference to the start of the allocated block and the collector would move the block to the free memory list. If the processor allocates memory, the block may end up being reallocated, destroying the current processing. This could result in a system failure.

A block of memory is implicitly available to be deallocated or returned to the list of free memory whenever there are no references to it. In a runtime environment supporting implicit memory management, a garbage collector usually scans or "walks" the dynamic memory from time to time looking for unreferenced blocks and returning them. The garbage collector starts at locations known to contain references to allocated

blocks. These locations are called "roots". The garbage collector examines the roots and when it finds a reference to an allocated block, it marks the block as referenced. If the block was unmarked, it recursively examines the block for references. When all the referenced blocks have been marked, a linear scan of all allocated memory is made and unreferenced blocks are swept into the free memory list. The memory may also be compacted by copying referenced blocks to lower memory locations that were occupied by unreferenced blocks and then updating references to point to the new locations for the allocated blocks.

The assumption that the garbage collector makes when attempting to scavenge or collect garbage is that all stacks are part of the root set of the walk. Thus, the stacks have to be fully described and walkable.

In programming environments like Smalltalk, where there are no type declarations, this is not particularly a problem. Only two different types of items, stack frames and objects, can be added to the stack. The garbage collector can easily distinguish between them and trace references relating to the objects.

However, the Java programming language also permits base types (i.e., integers) to be added to the stack. This greatly complicates matters because a stack walker has to be more aware how to view each stack slot. Base types slots must not be viewed as pointers (references), and must not be followed during a walk.

Further, the content of the stack may not be static, even during a single allocation. As a method runs, the stack is used as a temporary "scratch" space, and an integer might be pushed onto the stack or popped off it, or an object pushed or popped at any time. Therefore, it is important to know during the execution of a program that a particular memory location in the stack contains an integer or an object.

The changing content of a stack slot during method execution can be illustrated with the following simple bytecode sequence of the form:

```
ICONST 0
POP
NEW
POP
RETURN
```

As this is run, an integer, zero (0), is pushed onto the top of the stack, then popped so that the stack is empty. Then an object (pointer) is pushed onto the top of the stack, and then popped so that the stack is again empty. Schematically, the stack sequence is:

0
-
OBJECT
-

In this sequence, the constant 0 and the object share the same stack location as the program is running. Realistically, this sequence would never result in a garbage collection. However, in the naive case, if garbage collection did occur just after the integer was pushed onto the stack, the slot should be ignored, not walked, because it contains only an integer, whereas if a garbage collection occurred after the object had been pushed onto the stack, then the slot would have to be walked because it could contain the only reference to the object in the system. In addition, if the object on the stack had been moved to another location by compaction, then its pointer would have to be updated as well.

Thus, the stack walker has to have a scheme in place to determine which elements to walk and which to skip on the stack.

One solution proposed by Sun Microsystems, Inc in its U.S. Patent No. 5,668,999 for "System and Method for Pre-Verification of Stack Usage in Bytecode Program Loops", is to calculate the stack shapes for all bytecodes prior to program execution, and to store as a "snapshot", the state of a virtual stack paralleling typical stack operations required during the execution of a bytecode program. The virtual stack is used to verify that the stacks do not underflow or overflow. It

includes multiple, pre-set entry points, and can be used as a stack map in operations such as implicit memory management.

However, the creation of a virtual stack of the whole program can be costly in terms of processing time and memory allocation, when all that may be required is a stack mapping up to a specific program counter (PC) in the stack, for a garbage collector to operate a limited number of times during program execution.

## SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide mapping for any PC location on the stack. Then, if a garbage collection occurs, the shape of the stack can be determined for that part of the stack frame.

It is also an object of the present invention to provide a method for mapping the shape of a portion of the stack for use either statically, at method compilation, or dynamically, at runtime.

A further object of the invention is to provide memory optimizing stack mapping.

The stack mapper of the present invention seeks to determine the shape of the stack at a given PC. This is accomplished by locating all start points possible for a given method, that is, at all of the entry points for the method and all of the exception entry points, and trying to find a path from the beginning of the method to the PC in question. Once the path is found, a simulation is run

of the stack through that path, which is used as the virtual stack for the purposes of the garbage collector. Accordingly, the present invention provides a method for mapping a valid stack up to a destination program counter through mapping a path of control flow on the stack from any start point in a selected method to the destination program counter and simulating stack actions for executing bytecodes along said path. In order to map a path of control flow on the stack, bytecode sequences are processed linearly until the control flow is interrupted. As each bytecode sequence is processed, unprocessed targets from any branches in the sequence are recorded for future processing. The processing is repeatedly interactively, starting from the beginning of the method and then from each branch target until the destination program counter has been processed. Preferably a virtual stack is generated from the simulation, which is encoded and stored on either the stack or the heap.

## BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1A is a flow diagram outlining the steps taken by the stack mapper according to the present invention to map the shape of the stack to a given program counter during two passes;

Figure 1B is a flow diagram, similar to Figure 1A, illustrating the processing of a bytecode sequence at one point in the method of Figure 1A;

Figure 2 is a schematic diagram of a sample segment of stack slots for illustrating the method of operation of the invention;

Figure 3, consisting of Figures 3A through 3I, is a schematic illustration of the changes in three tables in memory used to track the processing of the individual program counters during the mapping of the sample segment of Figure 2, according to the preferred embodiment of the invention;

Figure 4 is a schematic illustration of a compiled method stored on the heap which includes static storage of a stack map generated during compilation of the method; and

Figure 5 is a schematic illustration of a stack constructed for a method which provides storage for a stack map generated dynamically at runtime.

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE INVENTION

"The Java Virtual Machine Specification" details the set of operations that a Java virtual machine must perform, and the associated stack actions. Not included in the Java specification are some more stringent requirements about code flow. These are specified in the bytecode verifier (discussed in detail in Sun's U.S. Patent No. 5,668,999, referenced above). Code sequences that allow for different stack shapes at a given PC are not allowed because they are not verifiable. Sequences that cause the stack to grow without bound are a good example.

Thus, the following code is not legal:

```
x:   ICONST1
     GOTO x
```

because it creates an infinite loop and a never-ending stack.

The present invention is described in the context of a Java programming environment. It can also apply to any environment that prohibits the use of illegal stack statements in a manner similar to that provided by the Java bytecode verifier.

The shape of the stack is determined by the control flows, the path or paths, within the method for which the stack frame was or will be constructed. Therefore, in the method of the present invention, a path from any start point of the method to a selected PC is located, and then the stack actions for the bytecodes along the path are simulated. The implementation of this method in the preferred embodiment is illustrated in more detail in the flow diagrams of Figures 1A and 1B, and discussed below.

Figure 2 is a sample of stack layout 200 for a method, to illustrate the preferred embodiment. (In the example, "JSR" refers to a jump to a subroutine, a branch with a return, and "IF EQ 0" is a comparison of the top of the stack against zero.) A linear scanning of these PCs as they are laid out in memory, starting at the beginning of the method and walking forward to a selected destination, such as PC 7, is not appropriate. The linear scan would omit the jump at PC 2 to the subroutine at PC 6,

resulting in a break in the stack model without knowledge of how to arrive at the selected PC.

Returning to Figure 1, the input to the method of the invention is the destination PC for the method and the storage area destination to which the resulting information on the stack shape will be written (block 100). When the mapping occurs at runtime, the definition of the storage destination will point to a location on the stack; when the mapping occurs at compile time, the pointer will be into an array for storage with the compiled method on the heap. The different uses of the invention for stack mapping at runtime and at compilation are discussed in greater detail below.

Memory for three tables, a seen list, a branch map table and a to be walked list, are allocated and the tables are initialized in memory (block 102). In the preferred embodiment, the memory requirement for the tables is sized in the following manner. For the seen list, one bit is reserved for each PC. This is determined by looking at the size of the bytecode array and reserving one bit for each bytecode. Similarly, two longs are allocated for each bytecode or PC in both the to be walked list and the branch map table. The bit vector format provides a fast implementation.

The three tables are illustrated schematically in Figure 3 for the code sequence given in Figure 2: Figure 3A shows the state of these tables at the beginning of the stack mapper's walk; Figures 3B through 3I show the

varying states of these tables as the stack mapper walks this code sequence.

The seen list is used in the first pass of the stack mapper to identify bytes which have already been walked, to avoid entering an infinite loop. At the beginning of the walk, no bytes in the given sequence are identified as having been seen. The to be walked list provides a list of all known entry points to the method. At the beginning of the stack mapper's walk, the to be walked list contains the entry point to the method at byte zero (0) and every exception handler address for the selected method. The branch map is initially empty.

Once these data structures are initialized, the first element from the to be walked list is selected (block 104) and the sequence of bytecodes is processed (block 106) in a straight line according to the following criteria or states and as illustrated in the flow diagram of Figure 1B. As each bytecode is selected for processing, it is added to the seen list (block 150). The actions taken in processing the bytecode are determined by the state that defines it:

state 0: flow unaffected (block 152), advance to next bytecode, if any (blocks 154, 156)

state 1: branch conditional (block 158), if branch target has not yet been seen (block 160), then add it to the to be walked list (block 162), and in any event, advance to next bytecode, if any (blocks 154, 156)

state 2: branch unconditional (block 164), if the branch target has not yet been seen (block 165), add it to the to be walked list (block 166) and end the straight walk (block 168)

state 3: jump to subroutine (JSR) (block 170), if branch target has not yet been seen (block 160), the add it to the to be walked list (block 162), and in any event, advance to the next bytecode, if any (blocks 154, 156)

state 4: return (block 172) ends the straight walk (block 168)

state 5: table bytecode (block 174), if branch targets have not yet been seen (block 165), then add them to the to be walked list (block 166), and end the straight walk (block 168)

state 6: wide bytecode (block 176), calculate size of bytecode to determine increment to next bytecode (block 178) and advance to next bytecode, if any (blocks 154, 156)

state 7: breakpoint bytecode (block 180), retrieve the actual bytecode and its state (block 182), and then process the actual bytecode (starting at block 150)

State 0 defines a byte that does not cause a branch or any control flow change. For example, in the sample

sequence of Figure 2, A LOAD does not affect the flow and would be processed as state 0.

A conditional branch (state 1) has two states; it can either fall through or go to destination. As the stack mapper processes a conditional branch, it assumes a fall through state, but adds the branch target to the to be walked list in order to process both sides of the branch. Figure 2 contains a conditional branch at bytes 4, 5. However, if a branch target has already been walked (according to the seen list), then the target is not added (block 156 in Figure 1B).

A JSR is a language construct used in languages like Java. It is similar to an unconditional branch, except that it includes a return, similar to a function call. It is treated in the same way as a conditional branch by the stack mapper. Figure 2 contains a JSR to byte 6 at byte 2.

Table bytecodes includes lookup tables and table switches (containing multiple comparisons and multiple branch targets). These are treated as an unconditional branch with multiple branches or targets; any targets not previously seen according to the seen list are added to the to be seen list.

Temporary betch and store instructions are normally one or two bytes long. One byte is for the bytecode and one byte is for the parameter unless it is inferred by the bytecode. However, Java includes an escape sequence which sets the parameters for the following bytecode as

larger than normal (wide bytecode). This affects the stack mapper only in how much the walk count is incremented for the next byte. It does not affect control.

Breakpoints are used for debugging purposes. The breakpoint has overlaid the actual bytecode in the sequence, so is replaced again by the actual bytecode. Processing of the bytecodes in the sequence continues until terminated (eg., by an unconditional branch or a return), or when there are no more bytecodes in the sequence. Returning to Figure 1A, if the selected PC was not seen during the walk because it is not found on the seen list (block 108), the next element on the to be walked list is selected (block 104) and the bytecode sequence from it processed (block 106) following the same steps in Figure 1B until the selected PC has been walked (block 108 in Figure 1A).

Thus, the processing of the bytecode sequence in Figure 2, given PC7 as the destination PC, would be performed as follows:

Figure 3A: At commencement, there would be only one element, PC 0 in the to be seen list 308.

Figure 3B: PC 0 is marked as "seen" in the seen list 310 and removed from the to be seen list 312. A LOAD does not affect the control flow; it is state 0. The stack mapper moves on to the next byte, PC 1.

Figure 3C: PC 1 is marked as "seen" in the seen list 314. The byte is again A LOAD, state 0, so the stack mapper moves on to the next PC.

Figure 3D: PC 2 ("JSR") is treated in the first pass as a conditional branch. Once PC 2 is added to the seen list 316, its target PC 6 is added to the to be walked list 320 and the branch PC 6 (destination PC 304) / PC 2 (source branch 306) is added to the branch list 318.

Figure 3E: The I LOAD of PC 3 is state 0, so the stack mapper moves to the next byte after adding PC 3 to the seen list 324.

Figure 3F: PC 4 is a conditional branch. After adding PC 4 to the seen list 326, the stack mapper attempts to add its target, PC 0, to the to be seen list 320, but cannot because PC 0 is already on the seen list 326.

Figure 3G: At the return of PC 5, code flow stops (state 4), ending the stack walk after PC 5 has been added to the seen list 328.

At this point, the stack mapper determines whether it has seen the destination PC 7 (as per block 108 in Figure 1A). Since it has not, the stack mapper begins processing a new line of bytecodes from the next entry on

the to be walked list (block 104). According to the sample of Figure 2, the next PC on the to be walked list 320 in Figure 3G) is PC 6, the conditional branch from PC 2. Therefore, after marking PC 6 as seen (seen list 330, Figure 3H), the stack mapper processed the PC according to state 0 and proceeds to the next bytecode, which is PC 7. PC 7 is marked as seen (seen list 332, Figure 3I), and the walk ends again because it has encountered a fresh return (state 4).

Once the selected PC has been walked (block 108), the path to the destination is calculated in reverse (block 110) by tracing from the destination PC 304 to the source PC 306 on the branch map list. In the example, the reverse flow is from PC 7 to PC 6 to PC 2. Because there is no comparable pairing of PC 2 with any other designated PC, it is assumed that PC 2 flows, in reverse, to PC 0. The reverse of this mapping provides the code flow from the beginning of the method to the destination PC 7, that is:

PC 0 -> PC 2 -> PC 6 -> PC 7.

This is the end of the first pass of the stack mapper over the bytecodes.

In the second pass, the stack mapper creates a simulation of the bytecodes (block 112) during which the stack mapper walks the path through the method determined from the first pass simulating what stack action(s) the virtual machine would perform for each object in this bytecode sequence. For many of the bytecode types (eg.,

A LOAD), the actions are table driven according to previously calculated stack action (pushes and pops) sequences.

Fifteen types of bytecodes are handled specially, mainly because instances of the same type may result in different stack action sequences (eg., different INVOKES may result in quite different work on the stack).

An appropriate table, listing the table-driven actions and the escape sequences in provided in the Appendix hereto. A virtual stack showing the stack shape up to the selected PC is constructed in memory previously allocated (block 114). In the preferred embodiment, one CPU word is used for each stack element.
The virtual stack is then recorded in a compressed encoded format that is readable by the virtual machine (block 116). In the preferred embodiment, each slot is compressed to a single bit that essentially distinguishes (for the use of the garbage collector) between objects and non-objects (eg., integers).

The compressed encoded stack map is stored statically in the compiled method or on the stack during dynamic mapping. In the case of static mapping, a stack map is generated and stored as the method is compiled on the heap. A typical compiled method shape for a Java method is illustrated schematically in Figure 4. The compiled method is made up of a number of fields, each four bytes in length, including the object header 400, bytecodes 402, start PC 404, class pointers 406, selector 408, Java flags 410 and laterals 414. According to the invention,

the compiled method also includes a field for the stack map 412. The stack map field 412 includes an array that encodes the information about the temps or local variables in the method generated by the stack mapper in the manner described above, and a linear stack map list that a garbage collector can use to access the stack shape for a given destination PC in the array by calculating the offset and locating the mapping bits in memory.

A stack map would normally be generated for static storage in the compiled method when the method includes an action that transfer control from that method, such as invokes, message sends, allocates and resolves.

The stack map can also be generated dynamically, for example, when an asynchronous event coincides with a garbage collection. To accommodate the map, in the preferred embodiment of the invention, empty storage is left on the stack.

Figure 5 illustrates a stack frame 500, having standard elements, such as an area for temps or arguments pushed by the method 502, laterals or the pointer the compiled method 504 (which also gives access to the stack map in the compiled method) and a back pointer 506 pointing to the previous stack frame. A small area of memory 508, possibly only four bytes, is left empty in the frame but tagged as needing dynamic mapping. An advantage of this is that if this stack frame 500 is deep in the stack, once the dynamic mapping has taken place, the frame will be undisturbed and is available for future activations.

The area on the stack for dynamic stack mapping 508 can be allocated whenever a special event occurs such as timer or asynchronous events and debugging, as well as for invokes, allocates and resolves discussed above.

While the invention has been particularly shown and described with respect to preferred embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and details may be made therein without departing from the spirit and scope of the invention.

6656550 092599

**APPENDIX**

Simulation Action Keys:

```
o.   pop.
u.   push int
U.   push object
d.   dup
1.   dupx1
2.   dupx2
3.   dup2
4.   dup2x1
5.   dup2x2
s.   swap
j.   jsr
m.   multianewarray
l.   ldc
i.   invoke(static|virtual|interface|special)
g.   get(field/static)
p.   put(field/static)
```

| # | Name | Simulation Action | Walk Action |
|---|---|---|---|
| 0 | nop | '' | 0x00 |
| 1 | aconstnull | 'U' | 0x00 |
| 2 | iconstm1 | 'u' | 0x00 |
| 3 | iconst0 | 'u' | 0x00 |
| 4 | iconst1 | 'u' | 0x00 |
| 5 | iconst2 | 'u' | 0x00 |
| 6 | iconst3 | 'u' | 0x00 |
| 7 | iconst4 | 'u' | 0x00 |
| 8 | iconst5 | 'u' | 0x00 |
| 9 | lconst0 | 'uu' | 0x00 |
| 10 | lconst1 | 'uu' | 0x00 |
| 11 | fconst0 | 'u' | 0x00 |
| 12 | fconst1 | 'u' | 0x00 |
| 13 | fconst2 | 'u' | 0x00 |
| 14 | dconst0 | 'uu' | 0x00 |
| 15 | dconst1 | 'uu' | 0x00 |
| 16 | bipush | 'u' | 0x00 |
| 17 | sipush | 'u' | 0x00 |
| 18 | ldc | 'l' | 0x00 |
| 19 | ldcw | 'l' | 0x00 |
| 20 | ldc2w | 'uu' | 0x00 |

| | | | |
|---|---|---|---|
| 21 | iload | 'u' | 0x00 |
| 22 | lload | 'uu' | 0x00 |
| 23 | fload | 'u' | 0x00 |
| 24 | dload | 'uu' | 0x00 |
| 25 | aload | 'U' | 0x00 |
| 26 | iload0 | 'u' | 0x00 |
| 27 | iload1 | 'u' | 0x00 |
| 28 | iload2 | 'u' | 0x00 |
| 29 | iload3 | 'u' | 0x00 |
| 30 | lload0 | 'uu' | 0x00 |
| 31 | lload1 | 'uu' | 0x00 |
| 32 | lload2 | 'uu' | 0x00 |
| 33 | lload3 | 'uu' | 0x00 |
| 34 | fload0 | 'u' | 0x00 |
| 35 | fload1 | 'u' | 0x00 |
| 36 | fload2 | 'u' | 0x00 |
| 37 | fload3 | 'u' | 0x00 |
| 38 | dload0 | 'uu' | 0x00 |
| 39 | dload1 | 'uu' | 0x00 |
| 40 | dload2 | 'uu' | 0x00 |
| 41 | dload3 | 'uu' | 0x00 |
| 42 | aload0 | 'U' | 0x00 |
| 43 | aload1 | 'U' | 0x00 |
| 44 | aload2 | 'U' | 0x00 |
| 45 | aload3 | 'U' | 0x00 |
| 46 | iaload | 'oou' | 0x00 |
| 47 | laload | 'oouu' | 0x00 |
| 48 | faload | 'oou' | 0x00 |
| 49 | daload | 'oouu' | 0x00 |
| 50 | aaload | 'ooU' | 0x00 |
| 51 | baload | 'oou' | 0x00 |
| 52 | caload | 'oou' | 0x00 |
| 53 | saload | 'oou' | 0x00 |
| 54 | istore | 'o' | 0x00 |
| 55 | lstore | 'oo' | 0x00 |
| 56 | fstore | 'o' | 0x00 |
| 57 | dstore | 'oo' | 0x00 |
| 58 | astore | 'o' | 0x00 |
| 59 | istore0 | 'o' | 0x00 |
| 60 | istore1 | 'o' | 0x00 |
| 61 | istore2 | 'o' | 0x00 |
| 62 | istore3 | 'o' | 0x00 |
| 63 | lstore0 | 'oo' | 0x00 |
| 64 | lstore1 | 'oo' | 0x00 |
| 65 | lstore2 | 'oo' | 0x00 |
| 66 | lstore3 | 'oo' | 0x00 |
| 67 | fstore0 | 'o' | 0x00 |

| | | | |
|---|---|---|---|
| 68 | fstore1 | 'o' | 0x00 |
| 69 | fstore2 | 'o' | 0x00 |
| 70 | fstore3 | 'o' | 0x00 |
| 71 | dstore0 | 'oo' | 0x00 |
| 72 | dstore1 | 'oo' | 0x00 |
| 73 | dstore2 | 'oo' | 0x00 |
| 74 | dstore3 | 'oo' | 0x00 |
| 75 | astore0 | 'o' | 0x00 |
| 76 | astore1 | 'o' | 0x00 |
| 77 | astore2 | 'o' | 0x00 |
| 78 | astore3 | 'o' | 0x00 |
| 79 | iastore | 'ooo' | 0x00 |
| 80 | lastore | 'oooo' | 0x00 |
| 81 | fastore | 'ooo' | 0x00 |
| 82 | dastore | 'oooo' | 0x00 |
| 83 | aastore | 'ooo' | 0x00 |
| 84 | bastore | 'ooo' | 0x00 |
| 85 | castore | 'ooo' | 0x00 |
| 86 | sastore | 'ooo' | 0x00 |
| 87 | pop | 'o' | 0x00 |
| 88 | pop2 | 'oo' | 0x00 |
| 89 | dup | 'd' | 0x00 |
| 90 | dupx1 | '1' | 0x00 |
| 91 | dupx2 | '2' | 0x00 |
| 92 | dup2 | '3' | 0x00 |
| 93 | dup2x1 | '4' | 0x00 |
| 94 | dup2x2 | '5' | 0x00 |
| 95 | swap | 's' | 0x00 |
| 96 | iadd | 'oou' | 0x00 |
| 97 | ladd | 'ooooou' | 0x00 |
| 98 | fadd | 'oou' | 0x00 |
| 99 | dadd | 'ooooou' | 0x00 |
| 100 | isub | 'oou' | 0x00 |
| 101 | lsub | 'ooooou' | 0x00 |
| 102 | fsub | 'oou' | 0x00 |
| 103 | dsub | 'ooooou' | 0x00 |
| 104 | imul | 'oou' | 0x00 |
| 105 | lmul | 'ooooou' | 0x00 |
| 106 | fmul | 'oou' | 0x00 |
| 107 | dmul | 'ooooou' | 0x00 |
| 108 | idiv | 'oou' | 0x00 |
| 109 | ldiv | 'ooooou' | 0x00 |
| 110 | fdiv | 'oou' | 0x00 |
| 111 | ddiv | 'ooooou' | 0x00 |
| 112 | irem | 'oou' | 0x00 |
| 113 | lrem | 'ooooou' | 0x00 |
| 114 | frem | 'oou' | 0x00 |

| | | | |
|---|---|---|---|
| 115 | drem | 'oooouu' | 0x00 |
| 116 | ineg | 'ou' | 0x00 |
| 117 | lneg | 'oouu' | 0x00 |
| 118 | fneg | 'ou' | 0x00 |
| 119 | dneg | 'oouu' | 0x00 |
| 120 | ishl | 'oou' | 0x00 |
| 121 | lshl | 'ooouu' | 0x00 |
| 122 | ishr | 'oou' | 0x00 |
| 123 | lshr | 'ooouu' | 0x00 |
| 124 | iushr | 'oou' | 0x00 |
| 125 | lushr | 'ooouu' | 0x00 |
| 126 | iand | 'oou' | 0x00 |
| 127 | land | 'oooouu' | 0x00 |
| 128 | ior | 'oou' | 0x00 |
| 129 | lor | 'oooouu' | 0x00 |
| 130 | ixor | 'oou' | 0x00 |
| 131 | lxor | 'oooouu' | 0x00 |
| 132 | iinc | '' | 0x00 |
| 133 | i2l | 'ouu' | 0x00 |
| 134 | i2f | 'ou' | 0x00 |
| 135 | i2d | 'ouu' | 0x00 |
| 136 | l2i | 'oou' | 0x00 |
| 137 | l2f | 'oou' | 0x00 |
| 138 | l2d | 'oouu' | 0x00 |
| 139 | f2i | 'ou' | 0x00 |
| 140 | f2l | 'ouu' | 0x00 |
| 141 | f2d | 'ouu' | 0x00 |
| 142 | d2i | 'oou' | 0x00 |
| 143 | d2l | 'oouu' | 0x00 |
| 144 | d2f | 'oou' | 0x00 |
| 145 | i2b | 'ou' | 0x00 |
| 146 | i2c | 'ou' | 0x00 |
| 147 | i2s | 'ou' | 0x00 |
| 148 | lcmp | 'oooou' | 0x00 |
| 149 | fcmpl | 'oou' | 0x00 |
| 150 | fcmpg | 'oou' | 0x00 |
| 151 | dcmpl | 'oooou' | 0x00 |
| 152 | dcmpg | 'oooou' | 0x00 |
| 153 | ifeq | 'o' | 0x01 |
| 154 | ifne | 'o' | 0x01 |
| 155 | iflt | 'o' | 0x01 |
| 156 | ifge | 'o' | 0x01 |
| 157 | ifgt | 'o' | 0x01 |
| 158 | ifle | 'o' | 0x01 |
| 159 | ificmpeq | 'oo' | 0x01 |
| 160 | ificmpne | 'oo' | 0x01 |
| 161 | ificmplt | 'oo' | 0x01 |

| | | | | |
|---|---|---|---|---|
| 162 | ificmpge | 'oo' | 0x01 |
| 163 | ificmpgt | 'oo' | 0x01 |
| 164 | ificmple | 'oo' | 0x01 |
| 165 | ifacmpeq | 'oo' | 0x01 |
| 166 | ifacmpne | 'oo' | 0x01 |
| 167 | goto | '' | 0x02 |
| 168 | jsr | 'j' | 0x03 |
| 169 | ret | 'on' | 0x04 |
| 170 | tableswitch | 'o' | 0x05 |
| 171 | lookupswitch | 'o' | 0x05 |
| 172 | ireturn | 'o' | 0x04 |
| 173 | lreturn | 'oo' | 0x04 |
| 174 | freturn | 'o' | 0x04 |
| 175 | dreturn | 'oo' | 0x04 |
| 176 | areturn | 'o' | 0x04 |
| 177 | return | '' | 0x04 |
| 178 | getstatic | 'g' | 0x00 |
| 179 | putstatic | 'p' | 0x00 |
| 180 | getfield | 'g' | 0x00 |
| 181 | putfield | 'p' | 0x00 |
| 182 | invokevirtual | 'i' | 0x00 |
| 183 | invokespecial | 'i' | 0x00 |
| 184 | invokestatic | 'i' | 0x00 |
| 185 | invokeinterface | 'i' | 0x00 |
| 187 | new | 'U' | 0x00 |
| 188 | newarray | 'oU' | 0x00 |
| 189 | anewarray | 'oU' | 0x00 |
| 190 | arraylength | 'ou' | 0x00 |
| 191 | athrow | 'ou' | 0x04 |
| 192 | checkcast | '' | 0x00 |
| 193 | instanceof | 'ou' | 0x00 |
| 194 | monitorenter | 'o' | 0x00 |
| 195 | monitorexit | 'o' | 0x00 |
| 196 | wide | '' | 0x06 |
| 197 | multianewarray | 'm' | 0x00 |
| 198 | ifnull | 'o' | 0x01 |
| 199 | ifnonnull | 'o' | 0x01 |
| 200 | gotow | '' | 0x02 |
| 201 | jsrw | 'j' | 0x03 |
| 202 | breakpoint | " | 0x07 |

# CLAIMS

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1. A method for mapping a valid stack up to a destination program counter, comprising:

mapping a path of control flow on the stack from any start point in a selected method to the destination program counter; and

simulating stack actions for executing bytecodes along said path.

2. The method of claim 1 wherein the step of mapping a path of control flow on the stack comprises:

processing a first linear bytecode sequence until the control flow is interrupted; and recording unprocessed targets from any branches in the first linear bytecode sequence for future processing.

3. The method of claim 2 wherein the step of mapping a path of control flow on the stack further comprises:

processing an additional bytecode linear sequence until the control flow is interrupted; and

recording unprocessed targets from any branches in the additional linear bytecode sequence for future processing, where the destination program counter was not reached during an earlier processing of a linear bytecode sequence.

4. The method of claim 2 wherein the step of processing any linear bytecode sequence comprises:

determining if a bytecode in said any linear bytecode sequence is a breakpoint with a pointer to bytecode data; and

replacing the breakpoint with the bytecode data.

5. The method of claim 3 wherein the step of processing any linear bytecode sequence comprises:

determining if a bytecode in said any linear bytecode sequence is a breakpoint with a pointer to bytecode data; and

replacing the breakpoint with the bytecode data.

6. The method of claim 1 wherein the step of simulating stack actions executing the bytecodes along the path further comprises generating a virtual stack.

7. The method of claim 6, further comprising:

encoding the virtual stack as a bitstring and storing the bitstring at a selected destination for use in memory management operations.

8. The method of claim 7, wherein the step of storing the bitstring comprises storing the bitstring to the selected method as compiled on a heap.

9. The method of claim 7, wherein the step of storing the bitstring comprises storing the bitstring to a pre-allocated area on the stack.

10. The method of claim 1 wherein the step of simulating stack actions executing the bytecodes along the path further comprises:

inserting pre-determined stack actions for bytecodes maintaining the control flow in the selected method; and

calculating stack actions for bytecodes transferring the control flow from the selected method.

11. A method for mapping a Java bytecode stack up to a destination program counter comprising:

mapping a path of control flow on the stack from any start point in a selected method to the destination program counter; and

simulating stack actions for executing bytecodes along said path.

12. The method of claim 11 wherein the step of mapping a path of control flow on the stack comprises:

processing a first linear bytecode sequence until the control flow is interrupted; and recording unprocessed targets from any branches in the first linear bytecode sequence for future processing.

13. The method of claim 12 wherein the step of mapping a path of control flow on the stack further comprises:

processing an additional bytecode linear sequence until the control flow is interrupted; and

recording unprocessed targets from any branches in the additional linear bytecode sequence for future processing, where the destination program counter was not reached during an earlier processing of a linear bytecode sequence.

14. The method of claim 12 wherein the step of processing any linear bytecode sequence comprises:

determining if a bytecode in said any linear bytecode sequence is a breakpoint with a pointer to bytecode data; and

replacing the breakpoint with the bytecode data.

15. The method of claim 13 wherein the step of processing any linear bytecode sequence comprises:

determining if a bytecode in said any linear bytecode sequence is a breakpoint with a pointer to bytecode data; and

replacing the breakpoint with the bytecode data.

16. The method of claim 11 wherein the step of simulating stack actions executing the bytecodes along the path further comprises generating a virtual stack.

17. The method of claim 16 further comprising:

encoding the virtual stack as a bitstring and storing the bitstring at a selected destination for use in memory management operations.

18. The method of claim 17, wherein the step of storing the bitstring comprises storing the bitstring to the selected method as compiled on a heap.

19. The method of claim 17, wherein the step of storing the bitstring comprises storing the bitstring to a pre-allocated area on the stack.

20. The method of claim 11 wherein the step of simulating stack actions executing the bytecodes along the path further comprises:

inserting pre-determined stack actions for bytecodes maintaining the control flow in the selected method; and

calculating stack actions for bytecodes transferring the control flow from the selected method.

21. A computer-readable memory for storing the instructions for use in the execution in a computer of the method of claim 1.

22. A computer readable memory for storing the instructions for use in the execution in a computer of the method of claim 11.

23. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform method steps for mapping a valid stack up to a destination program counter, said method steps comprising:

mapping a path of control flow on the stack from any start point in a selected method to the destination program counter; and

simulating stack actions for executing bytecodes along said path,

wherein the step of mapping a path of control flow on the stack comprises:

processing a first linear bytecode sequence until the control flow is interrupted; and

recording unprocessed targets from any branches in the first linear bytecode sequence for future processing, and

where the destination program counter was not reached during an earlier processing of a linear bytecode sequence,

processing an additional bytecode linear sequence until the control flow is interrupted; and

recording unprocessed targets from any branches in the additional linear bytecode sequence for future processing.

# MAPPING A STACK IN A STACK MACHINE ENVIRONMENT

## ABSTRACT OF THE DISCLOSURE

The stack mapper of the present invention seeks to determine the shape of the stack at a given program counter. This is accomplished by locating all start points possible for a given method, that is, at all of the entry points for the method and all of the exception entry points, and trying to find a path from the beginning of the method to the program counter in question. The mapper first tries to locate a linear path from the beginning of the method, and then interactively processes the sequence of bytes at each branch until the destination program counter is reached. Once the path is found, a simulation is run of the stack through that path, which is used as the virtual stack for the purposes of the garbage collector.

6004560 "6505705 "6505705 "6505705

FIGURE 1A

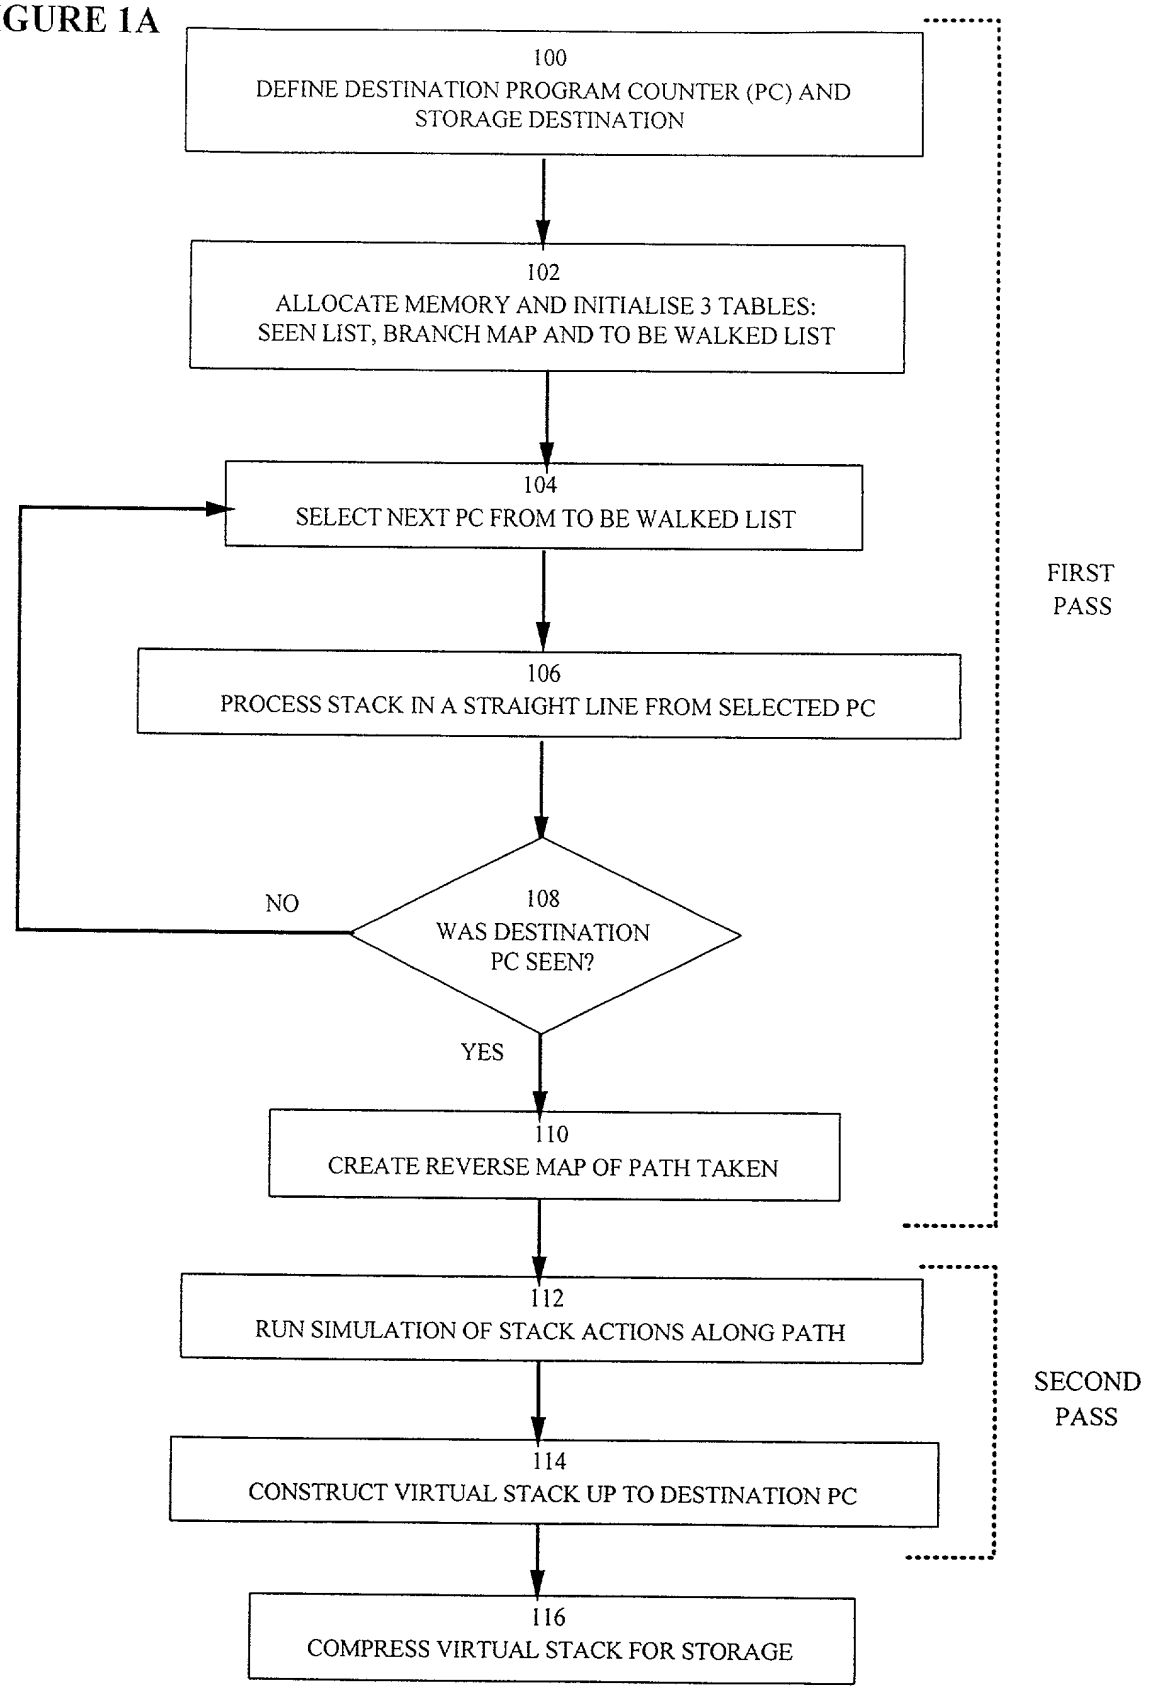


Docket No. 50277-2350

FIGURE 1B

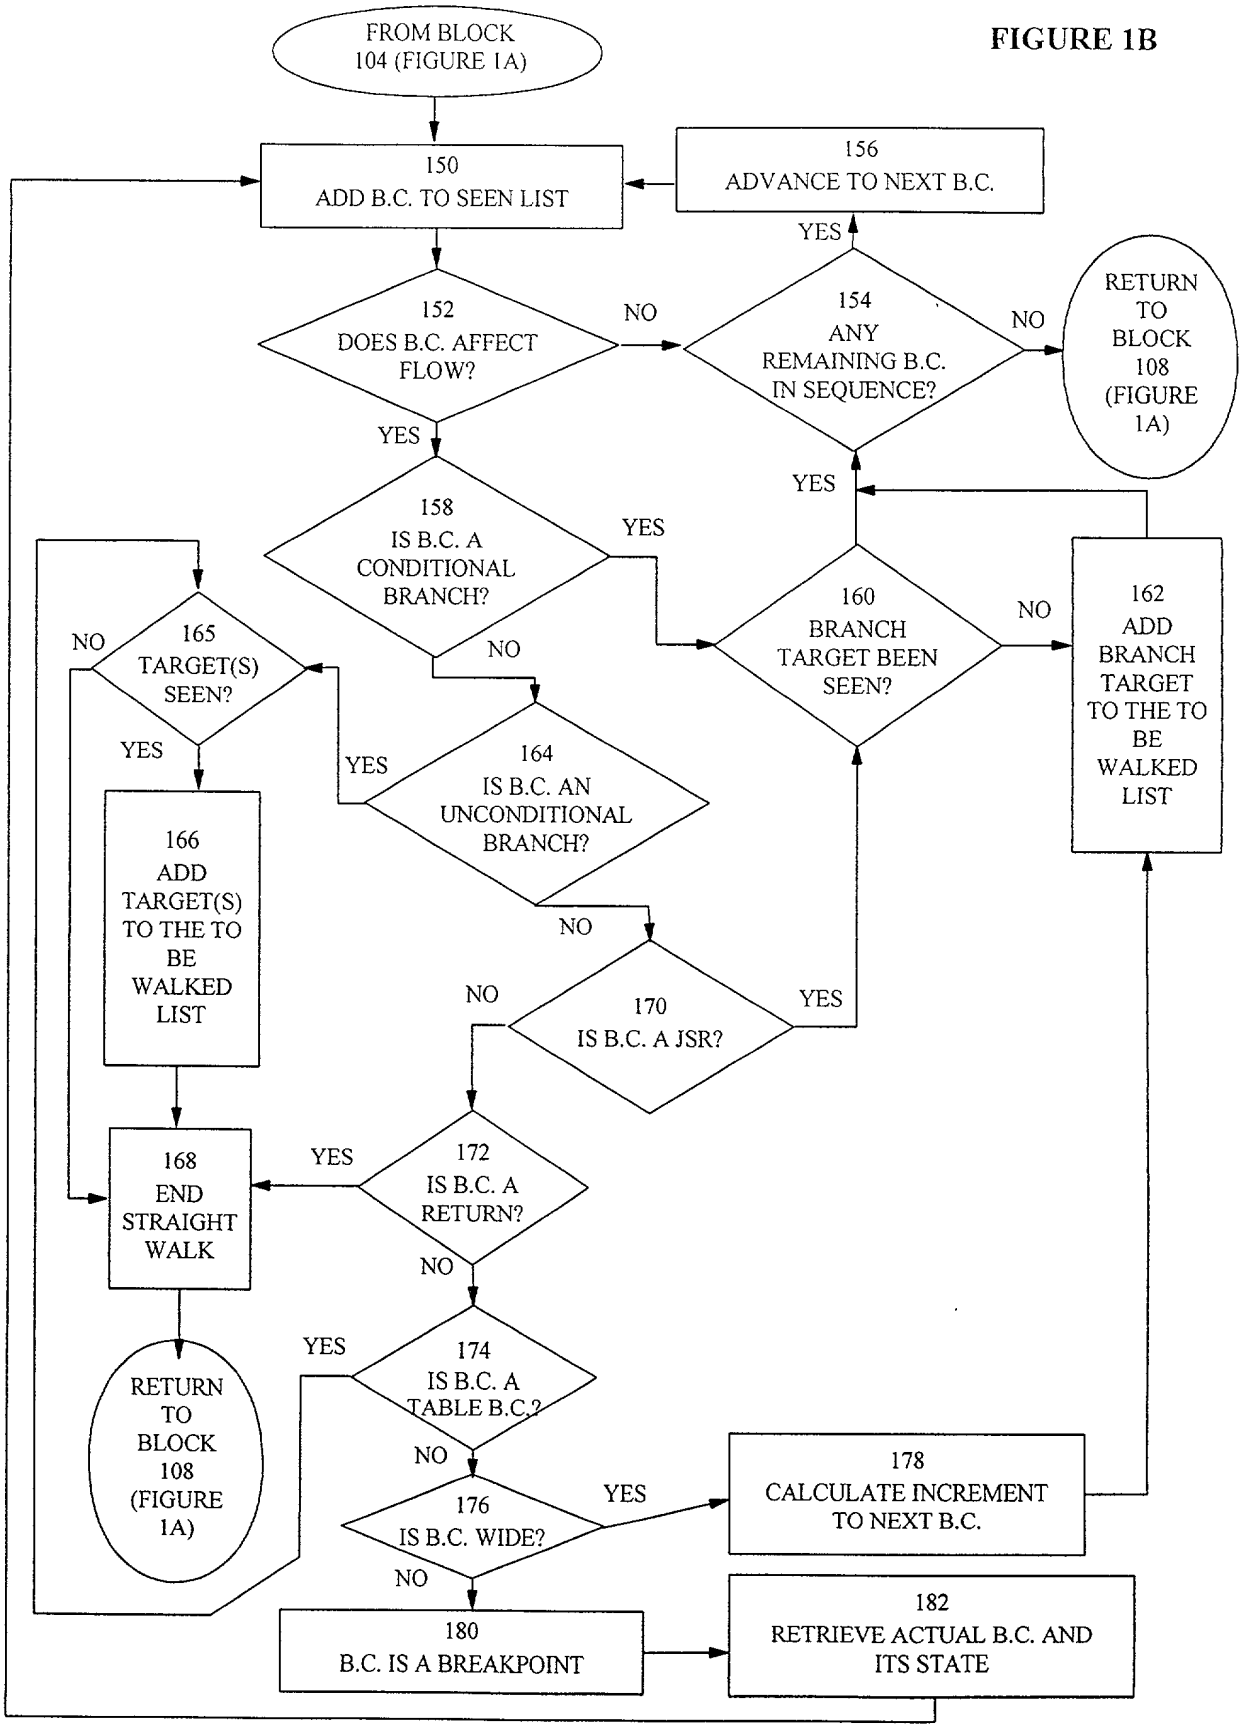
FROM BLOCK 104 (FIGURE 1A)
150 ADD B.C. TO SEEN LIST
156 ADVANCE TO NEXT B.C.
152 DOES B.C. AFFECT FLOW?
NO
154 ANY REMAINING B.C. IN SEQUENCE?
NO
RETURN TO BLOCK 108 (FIGURE 1A)
YES
YES
158 IS B.C. A CONDITIONAL BRANCH?
YES
160 BRANCH TARGET BEEN SEEN?
NO
162 ADD BRANCH TARGET TO THE TO BE WALKED LIST
NO
165 TARGET(S) SEEN?
YES
YES
164 IS B.C. AN UNCONDITIONAL BRANCH?
NO
166 ADD TARGET(S) TO THE TO BE WALKED LIST
NO
170 IS B.C. A JSR?
YES
NO
168 END STRAIGHT WALK
YES
172 IS B.C. A RETURN?
NO
RETURN TO BLOCK 108 (FIGURE 1A)
YES
174 IS B.C. A TABLE B.C.?
NO
176 IS B.C. WIDE?
YES
178 CALCULATE INCREMENT TO NEXT B.C.
NO
180 B.C. IS A BREAKPOINT
182 RETRIEVE ACTUAL B.C. AND ITS STATE

**FIGURE 2**

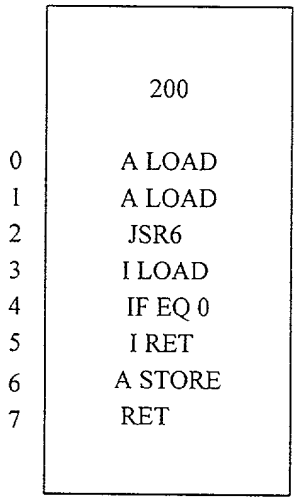


Sheet 2 of 5 U.S. Patent Dec. 5, 2000 6,157,988

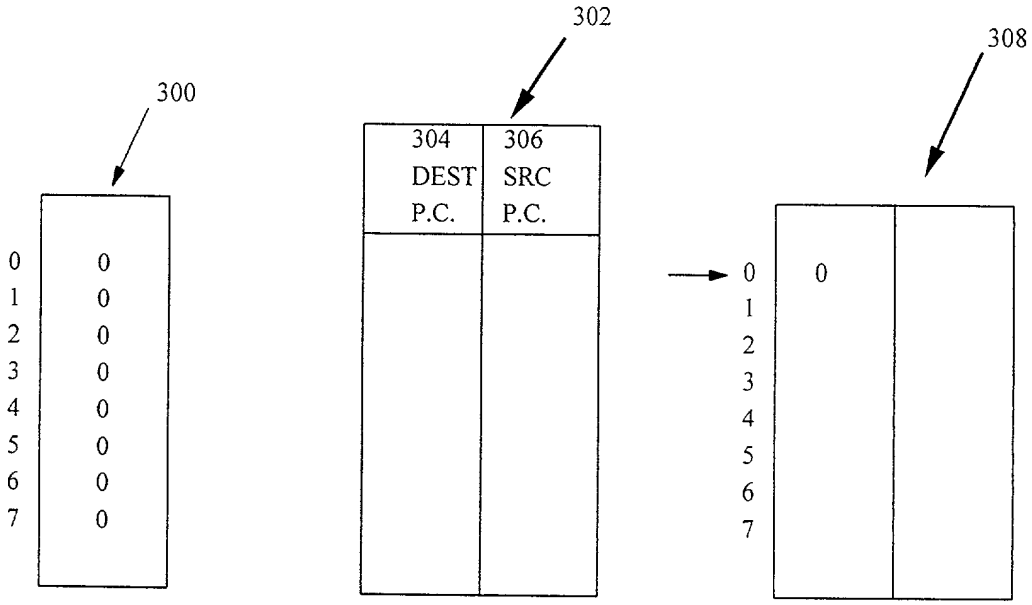
300
302
304
DEST
P.C.
306
SRC
P.C.
308
0
1
2
3
4
5
6
7
0
0
0
0
0
0
0
0
0
0
1
2
3
4
5
6
7
0

**FIGURE 3A**

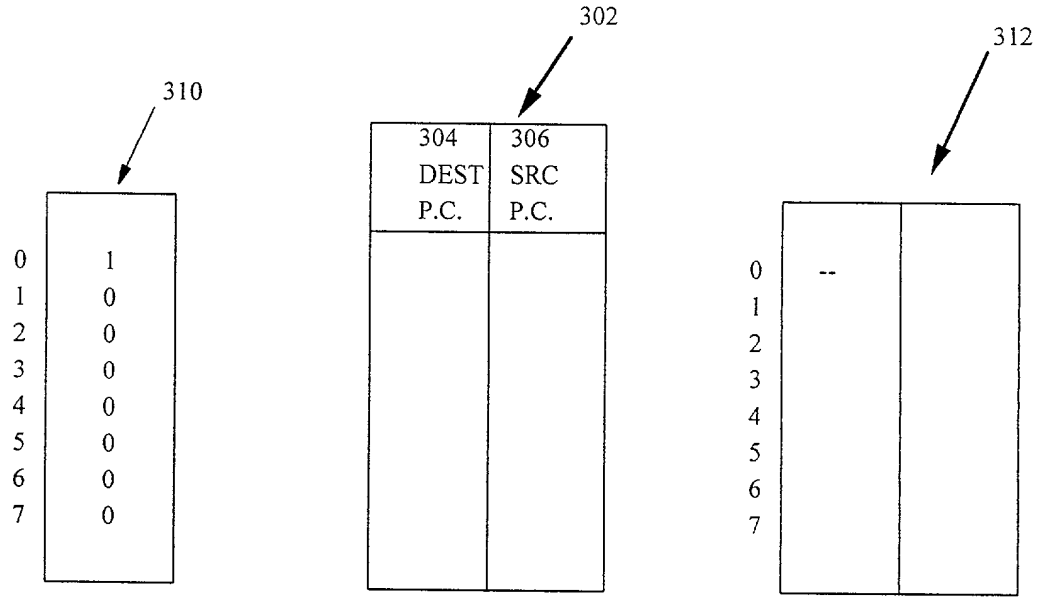
310
302
312
304
DEST
P.C.
306
SRC
P.C.
0 1
1 0
2 0
3 0
4 0
5 0
6 0
7 0
0 --
1
2
3
4
5
6
7

**FIGURE 3B**

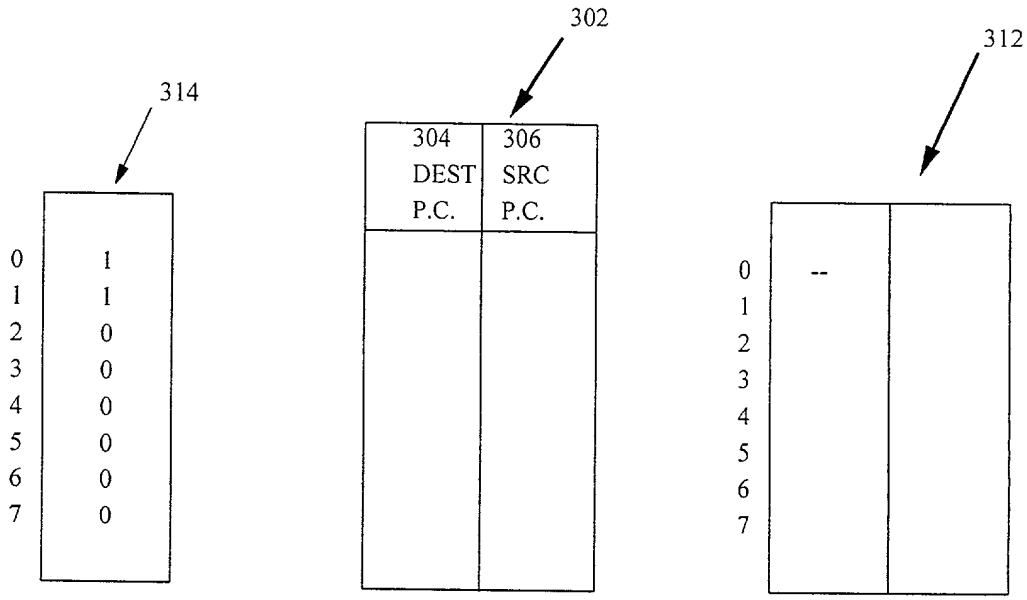
314
0 1
1 1
2 0
3 0
4 0
5 0
6 0
7 0
302
304 DEST P.C.
306 SRC P.C.
312
0 --
1
2
3
4
5
6
7

**FIGURE 3C**

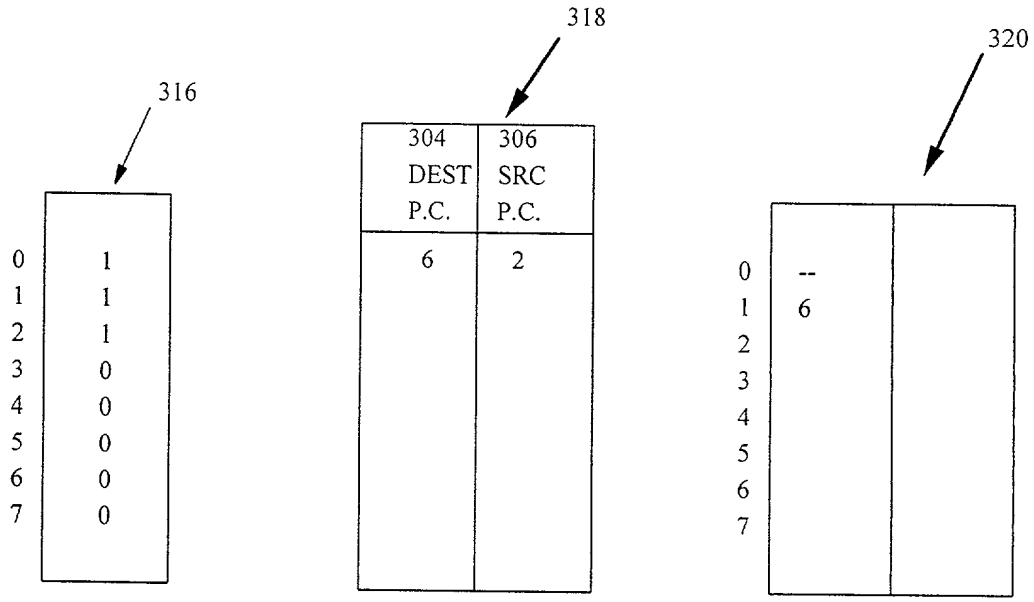


**FIGURE 3D**

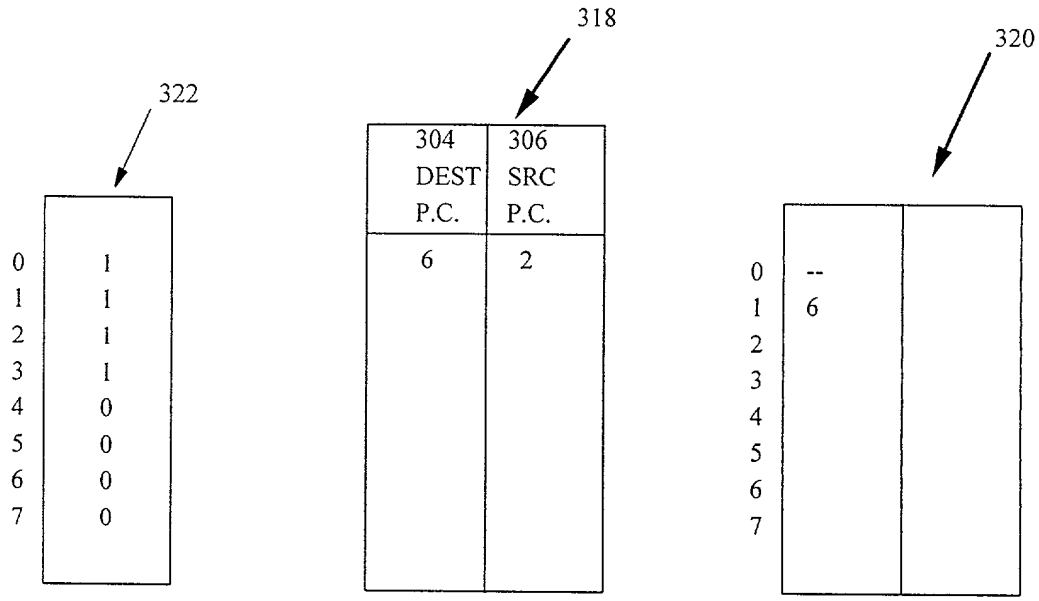
322
318
320
304
DEST
P.C.
306
SRC
P.C.
6
2
0 1
1 1
2 1
3 1
4 0
5 0
6 0
7 0
0 --
1 6
2
3
4
5
6
7

**FIGURE 3E**

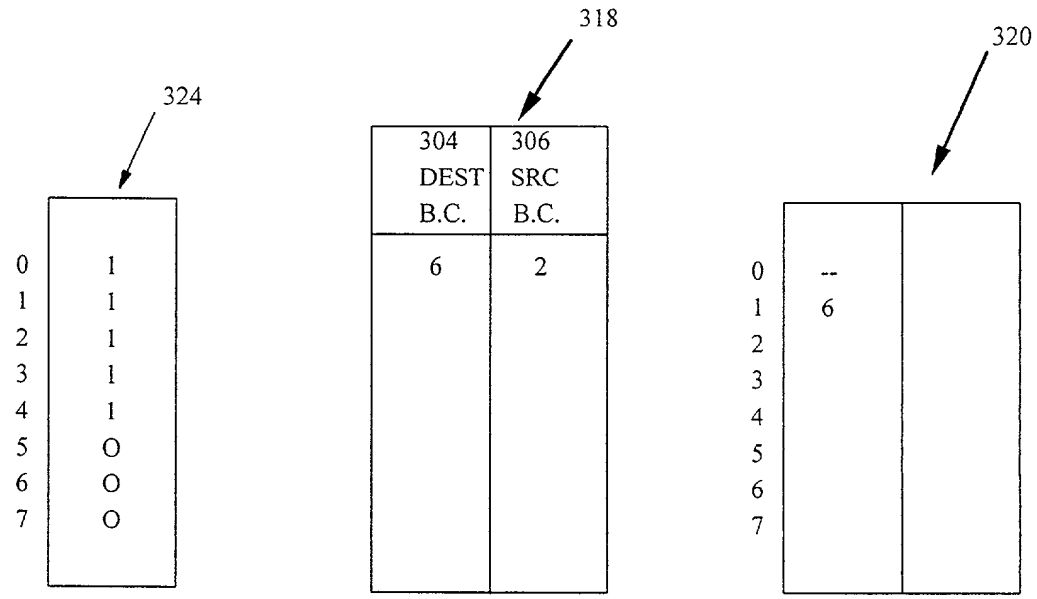


**FIGURE 3F**

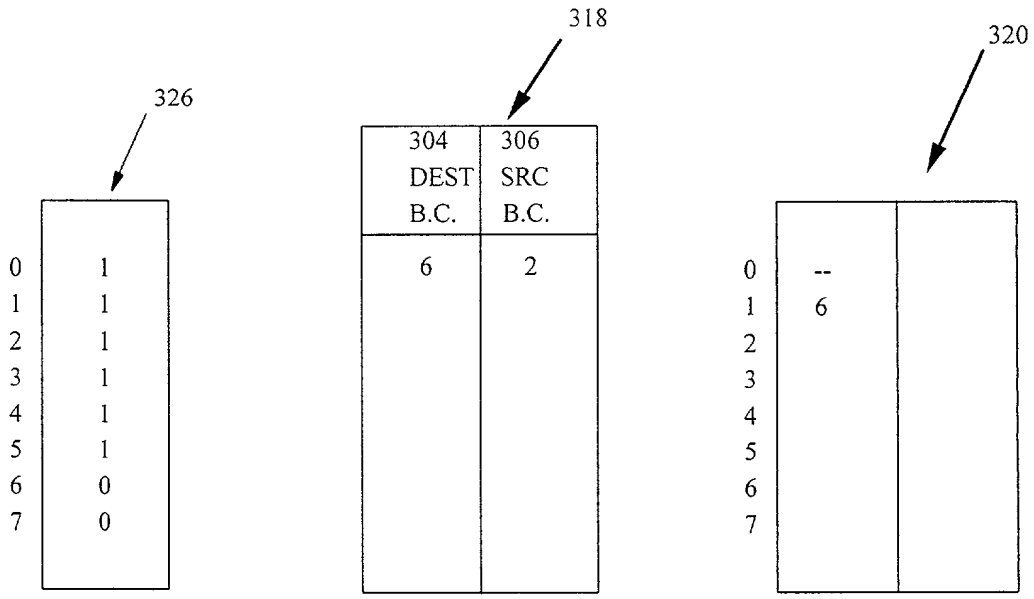


**FIGURE 3G**

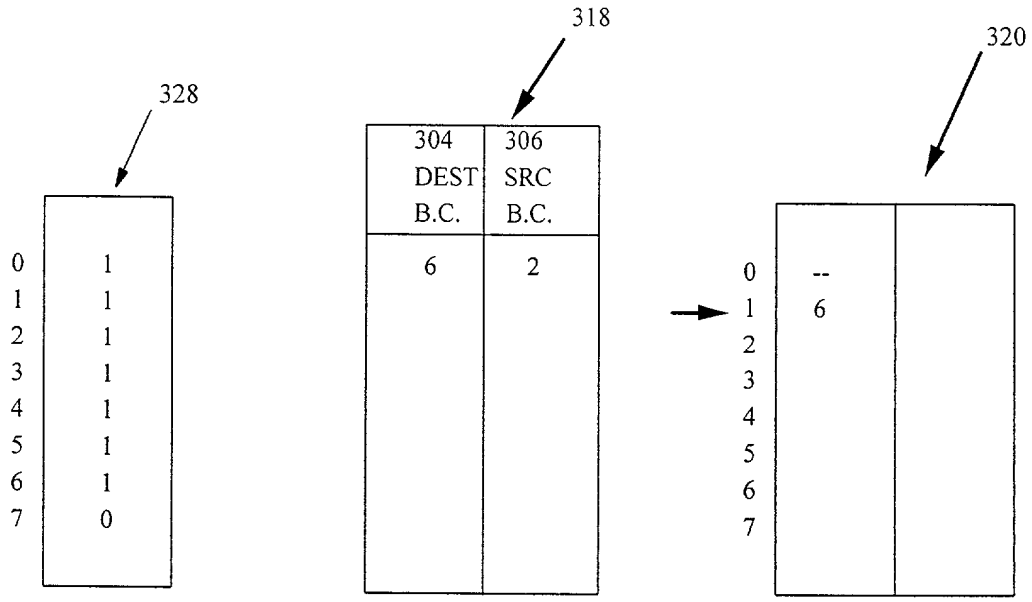
328
318
320
304
DEST
B.C.
306
SRC
B.C.
6
2
0 1
1 1
2 1
3 1
4 1
5 1
6 1
7 0
0 --
1 6
2
3
4
5
6
7

**FIGURE 3H**

Sheet 9 of 9 Attorney Docket No. 95-1-0

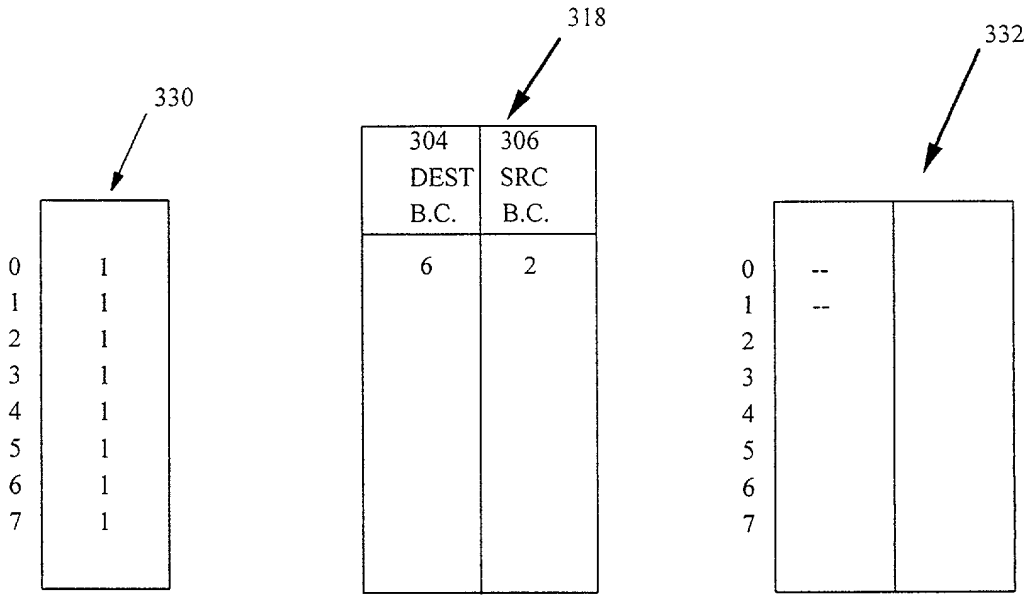


**FIGURE 3I**

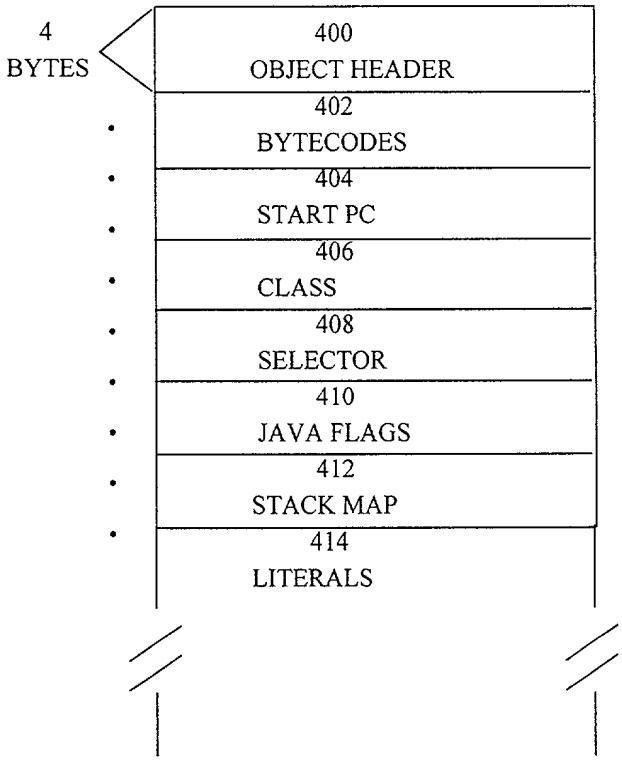
4
BYTES
400
OBJECT HEADER
402
BYTECODES
404
START PC
406
CLASS
408
SELECTOR
410
JAVA FLAGS
412
STACK MAP
414
LITERALS

**FIGURE 4**

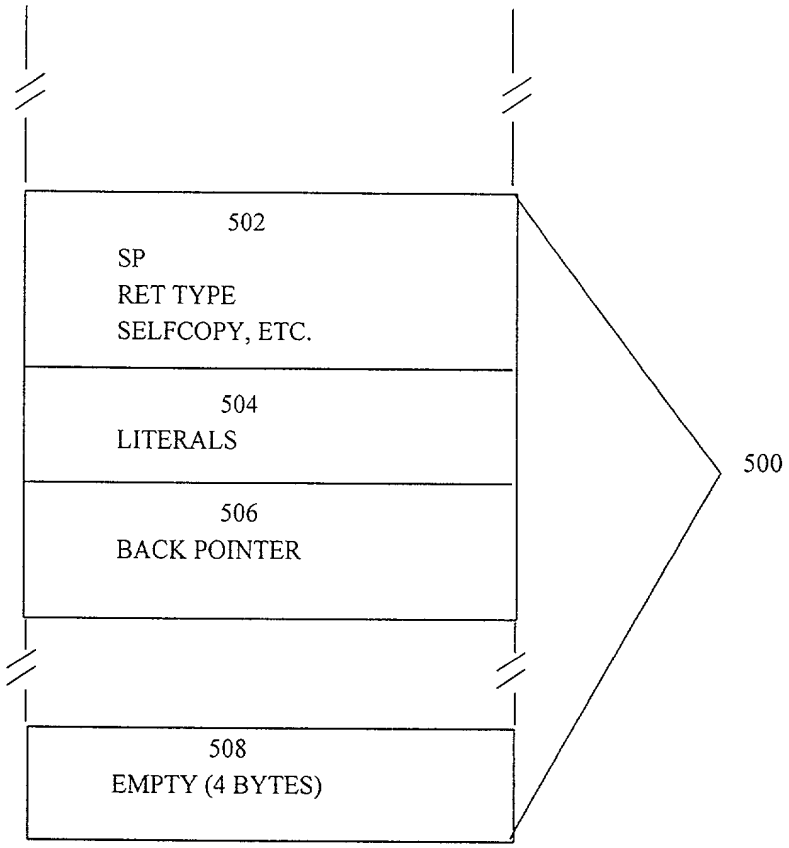
502
SP
RET TYPE
SELFCOPY, ETC.
504
LITERALS
506
BACK POINTER
508
EMPTY (4 BYTES)
500

**FIGURE 5**

SSM&P Docket No.: 12463
IBM Docket No.: CA998-012

# DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

**MAPPING A STACK IN A STACK MACHINE ENVIRONMENT**

the specification of which (check one)

X is attached hereto.

___ was filed on _____ as United States Application Number _____

or PCT International Application Number _____

and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) or §365(b) of any foreign application(s) for patent or inventor's certificate, or §365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application, having a filing date before that of the application on which priority is claimed:

| Prior Foreign Application(s) | | | Priority Claimed |
|---|---|---|---|
| 2,241,865 (Number) | Canada (Country) | 29 June 1998 (Day/Month/Year Filed) | X Yes ___ No |
| (Number) | (Country) | (Day/Month/Year Filed) | ___ Yes ___ No |
| (Number) | (Country) | (Day/Month/Year Filed) | ___ Yes ___ No |

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

| (Application Number) | (Filing Date) |
|---|---|
| (Application Number) | (Filing Date) |

I hereby claim the benefit under 35 U.S.C. §120 of any United States Application(s), or §365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States, or PCT International application in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in 37 CFR §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

| (Application Serial No.) | (Filing Date) | (Status) (patented, pending, abandoned) |
|---|---|---|
| (Application Serial No.) | (Filing Date) | (Status) (patented, pending, abandoned) |

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith (list name and registration number).

Manny W. Schecter (Reg. 31,722), Terry J. Ilardi (Reg. 29,936), Christopher A. Hughes (Reg. 26,914), Edward A. Pennington (Reg. 32,588), John E. Hoel (Reg. 26,279), Joseph C. Redmond, Jr. (Reg. 18,753), Douglas W. Cameron (Reg. No. 31,596), Kevin M. Jordan (Reg. No. 40,277), Stephen C. Kaufman (Reg. No. 29,551), Daniel P. Morris (Reg. No. 32,053), Paul J. Otterstedt (Reg. No. 37,411), Louis J. Percello (Reg. No. 33,206), Jay P. Sbrollini (Reg. No. 36,266), David M. Shofi (Reg. No. 39,835), Robert M. Trepp (Reg. No. 25,933) and Louis P. Herzberg (Reg. No. 41,500).

Send Correspondence to: Richard L. Catania, Scully, Scott, Murphy & Presser

400 Garden City Plaza, Garden City, New York 11530

Direct Telephone Calls to: (name and telephone number) Richard L. Catania, (516) 742-4343

Graham Chapman
**Full name of sole or first inventor**

Graham Chapman
Inventor's Signature

JUNE 1, 1999
Date

Nepean, Ontario, Canada
Residence

Canadian
Citizenship

30 Majestic Dr., Nepean, Ontario, Canada K2G 1C7
Post Office Address

SSM&P Docket No.: 12463
IBM Docket No.: CA998-012

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

Full name of second joint inventor, if any

John Duimovich — 1/June/99
Inventor's signature — Date

Ottawa, Ontario, Canada
Residence

Canadian
Citizenship

666 Rowanwood St., Ottawa, Ontario, Canada, K2A 3E4
Post Office Address

Trent Gray-Donald
Full name of third joint inventor, if any

1/June/99
Inventor's Signature — Date

Ottawa, Ontario, Canada
Residence

Canadian
Citizenship

11170 Cathcart Street, Ottawa, Ontario, Canada K1N 5B9
Post Office Address

Full name of fourth joint inventor, if any

Graeme Johnson — 1/June 99
Inventor's signature — Date

Ottawa, Ontario, Canada
Residence

Canadian
Citizenship

32 Fairbairn Street, Ottawa, Ontario, Canada K1S 1T3
Post Office Address

Andrew Low
Full name of fifth ~~joint~~ inventor, if any

1/June 99
Inventor's Signature — Date

OTTAWA
~~Stittsville~~, Ontario, Canada
Residence

Canadian
Citizenship

52 Stevenson Ave, OTTAWA K1Z 6N1
~~22 Elderwood Trail, Stittsville~~, Ontario, Canada ~~K2S 1T3~~
Post Office Address